

Modeling and Energy Optimization of LDPC Decoder Circuits with Timing Violations

François Leduc-Primeau, Frank R. Kschischang, and Warren J. Gross

Abstract

This paper proposes a “quasi-synchronous” design approach for signal processing circuits, in which timing violations are permitted, but without the need for a hardware compensation mechanism. The case of a low-density parity-check (LDPC) decoder is studied, and a method for accurately modeling the effect of timing violations at a high level of abstraction is presented. The error-correction performance of code ensembles is then evaluated using density evolution while taking into account the effect of timing faults. Following this, several quasi-synchronous LDPC decoder circuits based on the offset min-sum algorithm are optimized, providing a 23%–40% reduction in energy consumption or energy-delay product, while achieving the same performance and occupying the same area as conventional synchronous circuits.

I. INTRODUCTION

The time required for a signal to propagate through a CMOS circuit varies depending on several factors. Some of the variation results from physical limitations: the delay depends on the initial and final charge state of the circuit. Other variations are due to the difficulty (or impossibility) of controlling the fabrication process and the operating conditions of the circuit [1]. As process technologies approach atomic scales, the magnitude of these variations is increasing, and reducing the supply voltage to save energy increases the variations even further [2].

F. Leduc-Primeau is with the Electronics Department, Télécom Bretagne, Brest, France (e-mail: francois.leduc-primeau@telecom-bretagne.eu). W. J. Gross is with the Department of Electrical and Computer Engineering, McGill University, Montreal, Canada (e-mail: warren.gross@mcgill.ca). F. R. Kschischang is with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada (e-mail: frank@comm.utoronto.ca).

A preliminary version of this work was presented at the IEEE ICC 2015 conference.

The variation in propagation delay is a source of energy inefficiency for synchronous circuits since the clock period is determined by the worst delay. One approach to alleviate this problem is to allow timing violations to occur. While this would normally be catastrophic, some applications (in signal processing or in error-correcting decoding, for example) can tolerate a degradation in the operation of the circuit, either because an approximation to the ideal output suffices, or because the algorithm intrinsically rejects noise. This paper proposes an approach to the design of systems that are tolerant to timing violations. In particular we apply this approach to the design of energy-optimized low-density parity-check (LDPC) decoder circuits based on a state-of-the-art soft-input algorithm and architecture.

Other approaches have been previously proposed to build synchronous systems that can tolerate some timing violations. In *better than worst-case* (BTWC) [3] or *voltage over-scaled* (VOS) circuits, a mechanism is added to the circuit to compensate or recover from timing faults. One such method introduces special latches that can detect timing violations, and can trigger a restart of the computation when needed [4], [5]. Since the circuit's latency is increased significantly when a timing violation occurs, this approach is only suitable for tolerating small fault rates (e.g., 10^{-7}) and for applications where the circuit can be easily restarted, such as microprocessors that support speculative execution.

In most signal processing tasks, it is acceptable for the output to be non-deterministic, which creates more possibilities for dealing with timing violations. A seminal contribution in this area was the algorithmic noise tolerance (ANT) approach [6], [7], which is to allow timing violations to occur in the main processing block, while adding a separate reliable processing block with reduced precision that is used to bound the error of the main block, and provide algorithmic performance guarantees. The downside of the ANT approach is that it relies on the assumption that timing violations will first occur in the most significant bits. If that is not the case, the precision of the circuit can degrade to the precision of the auxiliary block, limiting the scheme's usefulness. For many circuits, including some adder circuits [8], this assumption does not hold. Furthermore, the addition of the reduced precision block and of a comparison circuit increases the area requirement.

We propose a design methodology for digital circuits with a relaxed synchronicity requirement that does not rely on any hardware compensation mechanism. Instead, we provide performance guarantees by re-analyzing the algorithm while taking into account the effect of timing violations.

We say that such systems are *quasi-synchronous*. LDPC decoding algorithms are good candidates for a quasi-synchronous implementation because their throughput and energy consumption are limiting factors in many applications, and like other signal processing algorithms, their performance is assessed in terms of expected values. Furthermore, since the algorithm is iterative, there is a possibility to optimize each iteration separately, and we show that this allows for additional energy savings.

The topic of unreliable LDPC decoders has been discussed in a number of contributions. Varshney studied the Gallager-A and the Sum-Product decoding algorithms when the computations and the message exchanges are “noisy”, and showed that the density evolution analysis still applies [9]. The Gallager-B algorithm was also analyzed under various scenarios [10]–[12]. A model for an unreliable quantized Min-Sum decoder was proposed in [13], which provided numerical evaluation of the density evolution equations as well as simulations of a finite-length decoder. Faulty finite-alphabet decoders were studied in [14], where it was proposed to model the decoder messages using conditional distributions that depend on the ideal messages. The quantized Min-Sum decoder was also analyzed in [15] for the case where faults are the result of storing decoder messages in an unreliable memory. The specific case of faults caused by delay variations in synchronous circuits is considered in [16], where a deviation model is proposed for binary-output circuits in which a deviation occurs probabilistically when the output of a circuit changes from one clock cycle to the next, but cannot occur if the output does not change. While none of these contributions explicitly consider the relationship between the reliability of the decoder’s implementation and the energy it consumes, there have been some recent developments in the analysis of the energy consumption of reliable decoders. Lower bounds for the scaling of the energy consumption of error-correction decoders in terms of the code length are derived in [17], and tighter lower bounds that apply to LDPC decoders are derived in [18]. The power required by regular LDPC decoders is also examined in [19], as part of the study of the total power required for transmitting and decoding the codewords.

In this paper, we present a modeling approach that provides an accurate representation of the deviations introduced in the output of an LDPC decoder processing circuit in the presence of occasional timing violations, while simultaneously measuring its energy consumption. We show that this model can be used as part of a density evolution analysis to evaluate the channel threshold and iterative performance of the decoder when affected by timing faults. Finally, we

show that under mild assumptions, the problem of minimizing the energy consumption of a quasi-synchronous decoder can be simplified to the energy minimization of a small test circuit, and present an approximate optimization method similar to Gear-Shift Decoding [20] that finds sequences of quasi-synchronous decoders that minimize decoding energy subject to performance constraints.

The remainder of the paper is organized as follows. Section II reviews LDPC codes and describes the circuit architecture of the decoder that is used to measure timing faults. Section III presents the *deviation* model that represents the effect of timing faults on the algorithm. Section IV then discusses the use of density evolution and of the deviation model to predict the performance of a decoder affected by timing faults. Finally, Section V presents the energy optimization strategy and results, and Section VI concludes the paper. Additional details on the CAD framework used for circuit measurements can be found in Appendix A, and Appendix B provides some details concerning the simulation of the test circuits.

II. LDPC DECODING ALGORITHM AND ARCHITECTURE

A. Code and Channel

We consider a communication scenario where a sequence of information bits is encoded using a binary LDPC code of length n . The LDPC code described by an $m \times n$ binary parity-check matrix $H = [h_{j,i}]$ consists of all length- n row vectors v satisfying the equation $vH^T = 0$. Equivalently, the code can be described by a bipartite Tanner graph with n *variable nodes* (VN) and m *check nodes* (CN) having an edge between the i -th variable node and the j -th check node if and only if $h_{j,i} \neq 0$. We assume that the LDPC code is regular, which means that in the code's Tanner graph each variable node has a fixed degree d_v and each check node has a fixed degree d_c .

Let us assume that the transmission takes place over the binary-input additive white Gaussian noise (BIAWGN) channel. A codeword $\mathbf{x} \in \{-1, 1\}^n$ is transmitted through the channel, which outputs the received vector $\mathbf{y} = \mathbf{x} + \mathbf{w}$, where \mathbf{w} is a vector of n independent and identically distributed (i.i.d.) zero-mean normal random variables with variance σ^2 . We use x_i and y_i to refer to the input and output of the channel at time i . The BIAWGN channel has the property of being output symmetric, meaning that $\phi_{y_i|x_i}(q | 1) = \phi_{y_i|x_i}(-q | -1)$, and memoryless, meaning that $\phi_{\mathbf{y}|\mathbf{x}}(\mathbf{q} | \mathbf{r}) = \prod_{i=1}^n \phi_{y_i|x_i}(q_i | r_i)$. Throughout the paper, $\phi(\cdot)$ denotes a probability density

function. The BIAWGN channel can also be described multiplicatively as $\mathbf{y} = \mathbf{x}\mathbf{z}$, where \mathbf{z} is a vector of i.i.d. normal random variables with mean 1 and variance σ^2 .

Let the *belief* output μ_i of the channel at time i be given by

$$\mu_i = \frac{\alpha y_i}{\sigma^2}, \quad (1)$$

with $\alpha > 0$. Note that if $\alpha = 2$ then μ_i is a log-likelihood ratio. Assuming that $x_i = 1$ was transmitted, then μ_i has a normal distribution with mean α/σ^2 and variance α^2/σ^2 . Writing $\rho = \alpha/\sigma^2$, we see that μ_i is Gaussian with mean ρ and variance $\alpha\rho$, that is, the distribution of μ_i is described by a single parameter ρ . We call this distribution a one-dimensional (1-D) normal distribution. The distribution of μ_i can also be specified using other equivalent parameters, such as the probability of error p_e , given by

$$p_e = \mathbb{P}(\mu_i < 0 | x_i = 1) = \mathbb{P}(\mu_i > 0 | x_i = -1) = \frac{1}{2} \operatorname{erfc}\left(\frac{1}{\sqrt{2\sigma^2}}\right) = \frac{1}{2} \operatorname{erfc}\left(\sqrt{\frac{\rho}{2\alpha}}\right), \quad (2)$$

where $\operatorname{erfc}(\cdot)$ is the complementary error function.

B. Decoding Algorithm

The well-known Offset Min-Sum (OMS) algorithm is a simplified version of the Sum-Product algorithm that can usually achieve similar error-correction performance. It has been widely used in implementations of LDPC decoders [21]–[23]. To make our decoder implementation more realistic and show the flexibility of our design framework, we present an algorithm and architecture that support a row-layered message-passing schedule. Architectures optimized for this schedule have proven effective for achieving efficient implementations of LDPC decoders [22]–[24]. Using a row-layered schedule also allows the decoder to be pipelined to increase the circuit's utilization. In a row-layered LDPC decoder, the rows of the parity-check matrix are partitioned into L sets called *layers*. To simplify the description of the decoding algorithm, we assume that all the columns in a given layer contain exactly one non-zero element. This implies that $L = d_v$. Note that codes with *at most* one non-zero element per column and per layer can also be supported by the same architecture, simply requiring a modification of the way algorithm variables are indexed.

Let us define a set \mathcal{L}_ℓ containing the indices of the rows of H that are part of layer ℓ , $\ell \in [1, L]$. We denote by $\mu_{i,j}^{(t)}$ a message sent from VN i to CN j during iteration t . and by $\lambda_{i,j}^{(t)}$ a message

sent from CN j to VN i . It is also useful to refer to the CN neighbor of a VN i that is part of layer ℓ . Because of the restriction mentioned above, there is exactly one such CN, and we denote its index by $J(i, \ell)$. Finally, we denote the channel information corresponding to the i -th codeword bit by $\mu_i^{(0)}$, since it also corresponds to the first message sent by a variable node i to all its neighboring check nodes.

The Offset Min-Sum algorithm used with a row-layered message-passing schedule is described in Algorithm 1. In the algorithm, $\mathcal{N}(j)$ denotes the set of indices corresponding to VNs that are neighbors of a check node j , and Λ_i represents the current sum of incoming messages at a VN i . The function $\min_{1,2}(S)$ returns the smallest and second smallest values in the set S , $C \geq 0$ is the offset parameter, and

$$\text{sgn}(x) = \begin{cases} 1 & \text{if } x \geq 0, \\ -1 & \text{if } x < 0. \end{cases}$$

C. Architecture

The Tanner graph of the code can also be used to represent the computations that must be performed by the decoder. At each decoding iteration, one message is sent from variable to check nodes on every edge of the graph, and again from check to variable nodes. We call a variable node processor (VNP) a circuit block that is responsible for generating messages sent by a variable node, and similarly a check node processor (CNP) a circuit block generating messages sent by a check node.

In a row-layered architecture in which the column weight of layer subsets is at most 1, there is at most one message to be sent and received for each variable node in a given layer. Therefore VNPs are responsible for sending and receiving one message per clock cycle. CNPs on the other hand receive and send d_c messages per clock cycle. At any given time, every VNP and CNP is mapped respectively to a VN and a CN in the Tanner graph. The routing of messages from VNPs to CNPs and back can be posed as two equivalent problems. One can fix the mapping of VNs to VNPs and of CNs to CNPs, and find a permutation of the message sequence that matches VNP outputs to CNP inputs, and another permutation that matches CNP outputs to VNP inputs. Alternatively, if VNPs process only one message at a time, one can fix the connections between VNPs and CNPs, and choose the assignment of VN to VNPs to achieve correct message routing.

Algorithm 1: OMS with a row-layered schedule.

input : $\{\mu_1^{(0)}, \mu_2^{(0)}, \dots, \mu_n^{(0)}\}$
output: $\{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n\}$

```

1 begin
    // Initialization
2    $\Lambda_i \leftarrow \mu_i^{(0)}, \forall i \in [1, n]$ 
3    $\lambda_{i,j}^{(0)} \leftarrow 0, \forall j \in [1, m], i \in \mathcal{N}(j)$ 
    // Decoding
4   for  $t \leftarrow 1$  to  $T$  do
5     for  $\ell \leftarrow 1$  to  $L$  do
        // VN to CN messages
6      $\mu_{i,J(i,\ell)}^{(t)} \leftarrow \Lambda_i - \lambda_{i,J(i,\ell)}^{(t-1)}, \forall i$ 
        // CN to VN messages
7     for  $j \in \mathcal{L}_\ell$  do
8        $[m_1, m_2] \leftarrow \min_{1,2}(\{|\mu_{i,j}^{(t)}| : i \in \mathcal{N}(j)\})$ 
9        $m_1 \leftarrow \max(0, m_1 - C)$ 
10       $m_2 \leftarrow \max(0, m_2 - C)$ 
11       $s_T \leftarrow \prod_{i \in \mathcal{N}(j)} \text{sgn}(\mu_{i,j}^{(t)})$ 
12      for  $i \in \mathcal{N}(j)$  do
13         $s_i \leftarrow s_T \cdot \text{sgn}(\mu_{i,j}^{(t)})$ 
14        if  $|\mu_{i,j}^{(t)}| = m_1$  then  $\lambda_{i,j}^{(t)} \leftarrow s_i \cdot m_2$  else  $\lambda_{i,j}^{(t)} \leftarrow s_i \cdot m_1$ 
        // VN update
15       $\Lambda_i \leftarrow \mu_{i,J(i,\ell)}^{(t)} + \lambda_{i,J(i,\ell)}^{(t)}, \forall i$ 
        // VN decision
16      for  $i \in \{1, 2, \dots, n\}$  do
17        if  $\Lambda_i > 0$  then  $\hat{x}_i \leftarrow 1$  else if  $\Lambda_i < 0$  then  $\hat{x}_i \leftarrow -1$  else
           $\hat{x}_i \leftarrow 1 \text{ or } -1 \text{ with equal probability}$ 

```

We choose the later approach because it allows studying the computation circuit without being concerned by the routing of messages.

The number of CNPs instantiated in the decoder can be adjusted based on throughput requirements from 1 to m/L (the number of rows in a layer). As the number of CNPs is varied, the number of VNPs will vary from d_c to n . An architecture diagram showing one VNP and one

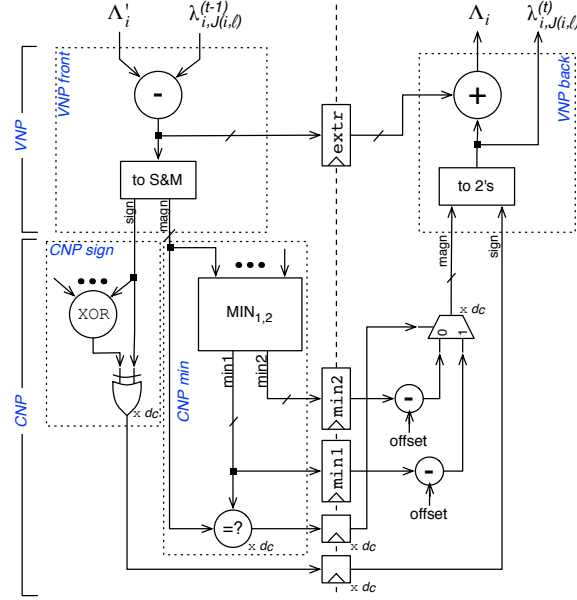


Fig. 1. Block diagram of the layered Offset Min-Sum decoder architecture.

CNP is shown in Fig. 1. In reality, a CNP is connected to $d_c - 1$ additional VNPs, which are not shown. The memories storing the belief totals Λ_i and the intrinsic beliefs $\lambda_{i,j}^{(t)}$ are also not shown. The part of the VNP responsible for sending a message to the CNP is called *VNP front* and the part responsible for processing a message received from a CNP is called the *VNP back*. The VNP front and back do not have to be simultaneously mapped to the same VN. This allows to easily vary the number of pipeline stages in the VNPs and CNPs. Fig. 1 shows the circuit with two pipeline stages.

Messages exchanged in the decoder are fixed-point numbers. The position of the binary point does not have an impact on the algorithm, and therefore the messages sent by VNs in the first iteration can be defined as rounding the result of (1) to the nearest integer, while choosing a suitable α . The number of bits in the quantization, the scaling factor α , and the OMS offset parameter are chosen based on a density evolution analysis of the algorithm (described in Section IV). We quantize decoder messages to 6 bits, which yields a decoder with approximately the same channel threshold as a floating-point decoder under a standard fault-free implementation.

In order to analyze a circuit that is representative of state-of-the-art architectures, we use an optimized architecture for finding the first two minima in each CNP. Our architecture is inspired

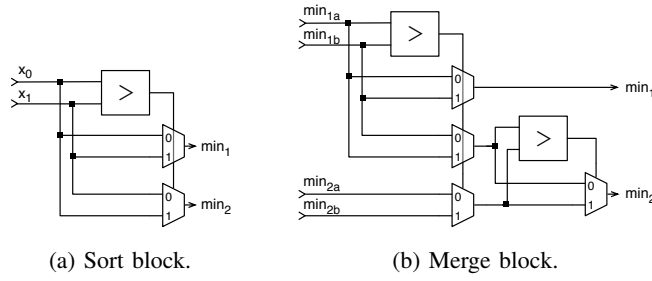


Fig. 2. Logic blocks used in the $\text{MIN}_{1,2}$ unit.

by the “tree structure” approach presented in [25], but requires fewer comparators. Each pair of CNP inputs is first sorted using the *Sort* block shown in Fig. 2a. These sorted pairs are then merged recursively using a tree of *Merge* blocks, shown in Fig. 2b. If the number of CNP inputs is odd, the input that cannot be paired is fed directly into a special merge block with 3 inputs, which can be obtained from the 4-input *Merge* block by removing the min_{2b} input and the bottom multiplexer.

Note that it is possible that changes to the architecture could increase or decrease the robustness of the decoder (see e.g. [26]), but this is outside the scope of this paper.

III. DEVIATION MODEL

A. Quasi-Synchronous Systems

We consider a synchronous system that permits timing violations without hardware compensation, resulting in what we call a *quasi-synchronous* system. Optimizing the energy consumption of these systems requires an accurate model of the impact of timing violations, and of the energy consumption. We propose to achieve this by characterizing a test circuit that is representative of the complete circuit implementation.

The term *deviation* refers to the effect of circuit faults on the result of a computation, and the deviation model is the bridge between the circuit characterization and the analysis of the algorithm. We reserve the term *error* for describing the algorithm, in the present case to refer to the incorrect detection of a transmitted symbol. A timing violation occurs in the circuit when the propagation delay between the input and output extends beyond a clock period. Modeling the deviations introduced by timing violations is challenging because they not only depend on the




Fig. 3. Computation tree of an LDPC decoder combined with the deviation model (for a regular LDPC code with $d_c = 4$ and $d_v = 3$).

current input to the circuit, but also on the state of the circuit before the new input was applied. In general, timing violations also depend on other dynamic factors and on process variations.

In this paper, we focus on the case where the output of the circuit is entirely determined by the current and previous inputs of the circuit, and by the nominal operating condition of the circuit. We denote by Γ the set of possible operating conditions, represented by vectors of parameters, and by $\gamma \in \Gamma$ a particular operating condition. For example, an operating condition might specify the supply voltage and clock period used in the circuit. We assume that all the parameters specified by γ are deterministic.

B. Test Circuit

The operation of an LDPC decoder can be represented using its one-iteration computation tree, which models the generation of a VN-to-CN message in terms of messages $\mu^{(t)}$ sent in the previous iteration. There are $(d_v - 1)$ check nodes in the tree. Each of these check nodes receives $(d_c - 1)$ messages from neighboring variable nodes, and generates a message sent to the one VN whose message was excluded from the computation. This VN then generates an extrinsic message based on the channel prior $\mu_i^{(0)}$ and on the messages received from neighboring check nodes. An example of a computation tree is shown within the dashed box in Fig. 3. For convenience, we choose to measure deviations on an implementation of this computation tree, so that measurements directly correspond with the progress made by the decoder in one iteration. As discussed in more details in Appendix B, the basic processing block of a row-layered decoder handles the messages to and from one check node. The test circuit is therefore built by re-using a basic block $d_v - 1$ times.

Since the test circuit is synchronous, we can represent it as a discrete-time system. Let X_k be the input at clock cycle k . When timing violations are allowed to occur, the corresponding¹ circuit output Z_k can be expressed as $Z_k = g(X_k, S_k)$, where S_k represents the state of the circuit at the beginning of cycle k , and g is some deterministic function. Equivalently, we can write $\mu_{i,j}^{(t+1)} = g(\boldsymbol{\mu}^{(t)}, S_k)$, where $\boldsymbol{\mu}^{(t)}$ is a vector containing all the VN-to-CN messages that form the input of the computation tree, and $\mu_{i,j}^{(t+1)}$ is a VN-to-CN message that will be sent in the next iteration.

A sequence of message vectors $\boldsymbol{\mu}^{(t)}$ can be mapped to a sequence of circuit inputs X_k in various ways. As is common with the type of decoder architecture considered here, we assume that all processing circuits are re-used several times during the same iteration t and layer ℓ . Therefore, for a fixed (t, ℓ) , the sequence of circuit inputs X_k forms an i.i.d. process. Since S_k depends on input X_{k-1} (and possibly also on other previous inputs), but not on X_k , S_k and X_k are independent. At the output, Z_k and Z_{k-1} are not independent, but it is possible to design the architecture so that correlated outputs are not associated with the same Tanner graph nodes or with neighboring nodes. This occurs naturally in a row-layered architecture, since each variable node is only updated once in each layer. Therefore, it is sufficient to consider the marginal distribution of the circuit's output, neglecting the correlation in successive outputs.

C. Deviation Model

We have seen above that a decoder message $\mu_{i,j}^{(t+1)}$ can be expressed as a function of the messages $\boldsymbol{\mu}^{(t)}$ received by the neighboring check nodes in the previous iteration and of the state of the processing circuit. To separate the deviations from the ideal operation of the decoder, it is helpful to decompose a decoding iteration into the ideal computation, followed by a transmission through a deviation channel. This model is shown in Fig. 3, where $\nu_{i,j}^{(t+1)}$ is the message that would be sent from variable node i to check node j during iteration $t + 1$ if no deviations had occurred during iteration t . For the first messages sent in the decoder at $t = 0$, the computation circuits are not used and therefore no deviation can occur, and we simply have $\mu_{i,j}^{(0)} = \nu_{i,j}^{(0)}$. Since we neglect correlations in successive circuit outputs, the deviation channel is memoryless.

¹The circuit could require one or several clock cycles to generate the first output, but this is irrelevant to the characterization of the computation.

Unlike typical channel models where the noise is independent from other variables in the system, the deviation $D_{i,j}^{(t)}$ is a function of the current circuit input $X_k = \boldsymbol{\mu}^{(t)}$ and of the current state S_k . However, modeling deviations directly in terms of $\boldsymbol{\mu}^{(t)}$ would make the model too complex, because of the large dimensionality of the input. To simplify the model, we consider only the value of the current output, and model deviations in terms of the conditional distribution $\phi(\mu_{i,j}^{(t+1)} \mid \nu_{i,j}^{(t+1)})$, an approach that was also used in [14]. To improve the accuracy of the model, it is also possible to consider the value of the transmitted bit x_i associated with VN i .

Since the faulty messages depend on the circuit state, the deviation model is obtained by averaging over the states S_k :

$$\phi(\mu_{i,j}^{(t+1)} \mid \nu_{i,j}^{(t+1)}, x_i) = \sum_{S_k} \phi(\mu_{i,j}^{(t+1)} \mid \nu_{i,j}^{(t+1)}, x_i, S_k) \phi(S_k), \quad (3)$$

which in practice can be done using a Monte-Carlo simulation of the test circuit.

D. Generalized Deviation Model

When evaluating deviations based on (3), it is important to keep in mind that S_k depends on previous circuit inputs. Under the assumption that the previous use of the circuit belonged to the same (t, ℓ) , $\phi(S_k)$ is a function of $\phi(\mu_{i,j}^{(t)})$. As a result, the model described by (3) is only valid for a fixed message distribution. Furthermore, because the message distribution depends on the transmitted codeword, the deviation model also depends on the transmitted codeword.

Let us first assume that the transmitted codeword is fixed. In this case, the message distribution $\phi(\mu_{i,j}^{(t)})$ depends on the channel noise, on the iteration index t , and on the operating condition of the circuit. Since the messages are affected by deviations for $t > 0$, only $\phi(\mu_{i,j}^{(0)})$ is known a priori. An obvious way to measure deviations is to perform a first evaluation of (3) using the known $\phi(\mu_{i,j}^{(0)})$, and to repeat the process for each subsequent decoding iteration. However, the resulting deviation model is of limited interest, since it depends on the specific message distributions in each iteration.

To generate a model that is independent of the iterative progress of the decoder, we first approximate $\phi(\mu_{i,j}^{(t)})$ as a 1-D Normal distribution with error rate parameter $p_e^{(t)}$ chosen such that

$$p_e^{(t)} = \mathbb{P}(x_i \mu_{i,j}^{(t)} < 0) + \frac{1}{2} \mathbb{P}(\mu_{i,j}^{(t)} = 0). \quad (4)$$

Note that while $\phi(\mu_{i,j}^{(0)})$ does correspond exactly to a 1-D Normal distribution, this is not necessarily the case after the first iteration. This approximation is the price to pay to obtain

a standalone deviation model, but note that exact message distributions can still be used when evaluating the performance of the faulty decoder. In fact, combining a density evolution based on exact distributions with a deviation model generated using 1-D Normal distributions leads to very accurate predictions in practice [27].

To construct the deviation model, we perform a number of Monte-Carlo simulations of (3) using 1-D input distributions with various $p_e^{(t)}$ values. Interpolation is then used to obtain a continuous model in $p_e^{(t)}$. The simulations are also performed for all operating conditions $\gamma \in \Gamma$. We therefore obtain a model that consists of a family of conditional distributions, indexed by $(p_e^{(t)}, \gamma)$, that we denote as

$$\phi^{(p_e^{(t)}, \gamma)} \left(\mu_{i,j}^{(t+1)} \mid \nu_{i,j}^{(t+1)}, x_i \right). \quad (5)$$

However, we generally omit the $(p_e^{(t)}, \gamma)$ superscript to simplify the notation. While measuring deviations, we also record the switching activity in the circuit, which is then used to construct an energy model that depends on γ and $p_e^{(t)}$, denoted as $c_\gamma(p_e^{(t)})$ (where c stands for “cost”).

To use the model, we first determine the error rate parameter $p_e^{(t)}$ corresponding to the distribution of the messages $\mu^{(t)}$ at the beginning of the iteration, and we then retrieve the appropriate conditional distribution, which also depends on the operating condition γ of the circuit. This conditional distribution then informs us of the statistics of deviations that occur at the end of the iteration, that is on messages sent in the next iteration.

As mentioned above, since $\phi(\mu_{i,j}^{(t)})$ depends on the transmitted codeword, this is also the case of $\phi(S_k)$ and of the deviation distributions. We show in Section IV that the codeword dependence is entirely contained within the deviation model and does not affect the analysis of the decoding performance, as long as the decoding algorithm and deviation model satisfy certain properties. Nonetheless, we would like to obtain a deviation model that does not depend on the transmitted codeword. This can be done when the objective is to predict the average performance of the decoder, rather than the performance for a particular codeword, since it is then sufficient to model the average behavior of the decoder. For the case where all codewords have an equal probability of being transmitted, we propose to perform the Monte-Carlo deviation measurements by randomly sampling transmitted codewords. This approach is supported by the experimental results presented in [27], which show that a deviation model constructed in this way can indeed accurately predict the average decoding performance.

IV. PERFORMANCE ANALYSIS

A. Standard Analysis Methods for LDPC Decoders

Density evolution (DE) is the most common tool used for predicting the error-correction performance of an LDPC decoder. The analysis relies on the assumption that messages passed in the Tanner graph are mutually independent, which holds as the code length goes to infinity [28]. Given the channel output probability distribution and the probability distribution of variable node to check node messages at the start of an iteration, DE computes the updated distribution of variable node to check node messages at the end of the decoding iteration. This computation can be performed iteratively to determine the message distribution after any number of decoding iterations. The validity of the analysis rests on two properties of the LDPC decoder. The first property is the conditional independence of errors, which states that the error-correction performance of the decoder is independent from the particular codeword that was transmitted. The second property states that the error-correction performance of a particular LDPC code concentrates around the performance measured on a cycle-free graph, as the code length goes to infinity.

Both properties were shown to hold in the context of reliable implementations [28]. It was also shown that the conditional independence of errors always holds when the channel is output symmetric and the decoder has a symmetry property. We can define a sufficient symmetry property of the decoder in terms of a message-update function $F_{i,j}$ that represents one complete iteration of the (ideal) decoding algorithm. Given a vector of all the messages $\boldsymbol{\mu}^{(t)}$ sent from variable nodes to check nodes at the start of iteration t and the channel information $\nu_i^{(0)}$ associated with variable node i , $F_{i,j}$ returns the next ideal message to be sent from a variable node i to a check node j : $\nu_{i,j}^{(t+1)} = F_{i,j}(\boldsymbol{\mu}^{(t)}, \nu_i^{(0)})$.

Definition 1. A message-update function $F_{i,j}$ is said to be symmetric with respect to a code C if

$$F_{i,j}(\boldsymbol{\mu}^{(t)}, \nu_i^{(0)}) = x_i F_{i,j}(\mathbf{x}\boldsymbol{\mu}^{(t)}, x_i \nu_i^{(0)})$$

for any $\boldsymbol{\mu}^{(t)}$, any $\nu_i^{(0)}$, and any codeword $\mathbf{x} \in C$.

In other words, a decoder's message-update function is symmetric if multiplying all the VN-to-CN belief messages sent at iteration t and the belief priors by a valid codeword $\mathbf{x} \in C$ is

equivalent to multiplying the next messages sent at iteration $t + 1$ by that same codeword. Note that the symmetry condition in Definition 1 is implied by the check node and variable node symmetry conditions in [28, Def. 1].

B. Applicability of Density Evolution

In order to use density evolution to predict the performance of long finite-length codes, the decoder must satisfy the two properties stated in Section IV-A, namely the conditional independence of errors and the convergence to the cycle-free case. We first present some properties of the decoding algorithm and of the deviation model that are sufficient to ensure the conditional independence of errors.

Using the multiplicative description of the BIAWGN channel, the vector received by the decoder is given by $\mathbf{y} = \mathbf{x}\mathbf{z}$ when a codeword \mathbf{x} is transmitted, or by $\mathbf{y} = \mathbf{z}$ when the all-one codeword is transmitted. In a reliable decoder, messages are completely determined by the received vector, but in a faulty decoder, there is additional randomness that results from the deviations. Therefore, we represent messages in terms of conditional probability distributions given $\mathbf{x}\mathbf{z}$. Since we are concerned with a fixed-point circuit implementation of the decoder, we can assume that messages are integers from the set $\{-Q, -Q + 1, \dots, Q\}$, where $Q > 0$ is the largest message magnitude that can be represented.

Definition 2. We say that a message distribution $\phi_{\mu_{i,j}|\mathbf{y}}(\mu|\mathbf{x}\mathbf{z})$ is symmetric if

$$\phi_{\mu_{i,j}|\mathbf{y}}(\mu|\mathbf{x}\mathbf{z}) = \phi_{\mu_{i,j}|\mathbf{y}}(x_i\mu|\mathbf{z}).$$

If a message has a symmetric distribution, its error probability as defined in (4) is the same whether $\mathbf{x}\mathbf{z}$ or \mathbf{z} is received. Similarly to the results presented in [14], we can show that the symmetry of message distributions is preserved when the message-update function is symmetric.

Lemma 1. If $F_{i,j}$ is a symmetric message-update function and if $\mu_i^{(0)}$ and $\mu_{i,j}^{(t)}$ have symmetric distributions for all (i, j) , the next ideal messages $\nu_{i,j}^{(t+1)}$ also have symmetric distributions.

Proof: We can express the distribution of the next ideal message from VN i to CN j as

$$\phi_{\nu_{i,j}^{(t+1)}|\mathbf{y}}(\nu|\mathbf{x}\mathbf{z}) = \sum_{(\boldsymbol{\mu}, \mu_i^{(0)}) \in R} \phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\boldsymbol{\mu}|\mathbf{x}\mathbf{z}) \phi_{\mu_i^{(0)}|\mathbf{y}}(\mu_i^{(0)}|\mathbf{x}\mathbf{z}), \quad (6)$$

where $R = \{(\boldsymbol{\mu}, \mu_i^{(0)}) : F_{i,j}(\boldsymbol{\mu}, \mu_i^{(0)}) = \nu\}$.

Assuming that the elements of the VN-to-CN message vector $\boldsymbol{\mu}^{(t)}$ are independent and that each $\mu_{i,j}^{(t)}$ has a symmetric distribution,

$$\phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\boldsymbol{\mu} | \mathbf{x}\mathbf{z}) = \prod_k \phi_{\mu_k^{(t)}|\mathbf{y}}(\mu_k | \mathbf{x}\mathbf{z}) = \prod_k \phi_{\mu_k^{(t)}|\mathbf{y}}(x_k \mu_k | \mathbf{z}) = \phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\mathbf{x}\boldsymbol{\mu} | \mathbf{z}),$$

and since the channel output $\mu_i^{(0)}$ also has a symmetric distribution,

$$\phi_{\mu_i^{(0)}|\mathbf{y}}(\mu_i^{(0)} | \mathbf{x}\mathbf{z}) = \phi_{\mu_i^{(0)}|\mathbf{y}}(x_i \mu_i^{(0)} | \mathbf{z}).$$

Therefore, we can rewrite (6) as

$$\phi_{\nu_{i,j}^{(t+1)}|\mathbf{y}}(\nu | \mathbf{x}\mathbf{z}) = \sum_{(\boldsymbol{\mu}, \mu_i^{(0)}) \in R} \phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\mathbf{x}\boldsymbol{\mu} | \mathbf{z}) \phi_{\mu_i^{(0)}|\mathbf{y}}(x_i \mu_i^{(0)} | \mathbf{z}). \quad (7)$$

Finally, letting $\boldsymbol{\mu}' = \mathbf{x}\boldsymbol{\mu}^{(t)}$ and $\nu'_i = x_i \mu_i^{(0)}$, (7) becomes

$$\phi_{\nu_{i,j}^{(t+1)}|\mathbf{y}}(\nu | \mathbf{x}\mathbf{z}) = \sum_{(\boldsymbol{\mu}', \nu'_i) \in R'} \phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\boldsymbol{\mu}' | \mathbf{z}) \phi_{\mu_i^{(0)}|\mathbf{y}}(\nu'_i | \mathbf{z}),$$

where $R' = \{(\boldsymbol{\mu}', \nu'_i) : F_{i,j}(\mathbf{x}\boldsymbol{\mu}', x_i \nu'_i) = \nu\}$. Since $F_{i,j}$ is symmetric, we can also express R' as

$$R' = \{(\boldsymbol{\mu}', \nu'_i) : F_{i,j}(\boldsymbol{\mu}', \nu'_i) = x_i \nu\},$$

and therefore,

$$\begin{aligned} \phi_{\nu_{i,j}^{(t+1)}|\mathbf{y}}(x_i \nu | \mathbf{z}) &= \sum_{(\boldsymbol{\mu}', \nu'_i) \in R'} \phi_{\boldsymbol{\mu}^{(t)}|\mathbf{y}}(\boldsymbol{\mu}' | \mathbf{z}) \phi_{\mu_i^{(0)}|\mathbf{y}}(\nu'_i | \mathbf{z}) \\ &= \phi_{\nu_{i,j}^{(t+1)}|\mathbf{y}}(\nu | \mathbf{x}\mathbf{z}), \end{aligned}$$

indicating that the next ideal messages have symmetric distributions. ■

To establish the conditional independence of errors under the proposed deviation model, we first define some properties of the deviation.

Definition 3. We say that the deviation model is symmetric if

$$\phi_{\mu_{i,j}^{(t)}|\nu_{i,j}^{(t)},\mathbf{y}}(\mu | \nu, \mathbf{x}\mathbf{z}) = \phi_{\mu_{i,j}^{(t)}|\nu_{i,j}^{(t)},\mathbf{y}}(\mu | \nu, \mathbf{z}) = \phi_{\mu_{i,j}^{(t)}|\nu_{i,j}^{(t)},\mathbf{y}}(-\mu | -\nu, \mathbf{z}).$$

Definition 4. We say that the deviation model is weakly symmetric (WS) if

$$\phi_{\mu_{i,j}^{(t)}|\nu_{i,j}^{(t)},\mathbf{y}}(\mu | \nu, \mathbf{x}\mathbf{z}) = \phi_{\mu_{i,j}^{(t)}|\nu_{i,j}^{(t)},\mathbf{y}}(x_i \mu | x_i \nu, \mathbf{z}).$$

Note that if the model satisfies the symmetry condition, it also satisfies the weak symmetry condition, since $x_i \in \{-1, 1\}$. We then have the following Lemma.

Lemma 2. *If a decoder having a symmetric message-update function and taking its inputs from an output-symmetric communication channel is affected by weakly symmetric deviations, its message error probability at any iteration $t \geq 0$ is independent of the transmitted codeword.*

Proof: Similarly to the approach used in [29, Lemma 4.90] and [9], we want to show that the probability that messages are in error is the same whether \mathbf{xz} or \mathbf{z} is received. This is the case if the faulty messages $\mu_{i,j}^{(t)}$ have a symmetric distribution for all $t \geq 0$ and all (i, j) .

Since the communication channel is output symmetric and since no deviations can occur before the first iteration, messages $\mu_{i,j}^{(0)} = \nu_{i,j}^{(0)}$ have a symmetric distribution. We proceed by induction to establish the symmetry of the messages for $t > 0$. We start by assuming that

$$\phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(\nu | \mathbf{xz}) = \phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(x_i \nu | \mathbf{z}) \quad (8)$$

also holds for $t > 0$.

Using Definition 4 and (8), we can write the faulty message distribution as

$$\begin{aligned} \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(\mu | \mathbf{xz}) &= \sum_{\nu=-Q}^Q \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(\mu | \nu, \mathbf{xz}) \phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(\nu | \mathbf{xz}) \\ &= \sum_{\nu=-Q}^Q \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(x_i \mu | x_i \nu, \mathbf{z}) \phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(x_i \nu | \mathbf{z}) \\ &= \sum_{\nu'=-x_i Q}^{x_i Q} \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(x_i \mu | \nu', \mathbf{z}) \phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(\nu' | \mathbf{z}) \\ &= \sum_{\nu'=-Q}^Q \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(x_i \mu | \nu', \mathbf{z}) \phi_{\nu_{i,j}^{(t)}|\mathbf{y}}(\nu' | \mathbf{z}) \\ &= \phi_{\mu_{i,j}^{(t)}|\mathbf{y}}(x_i \mu | \mathbf{z}). \end{aligned}$$

where the third equality is obtained using the substitution $\nu' = x_i \nu$. We conclude that the faulty messages have a symmetric distribution. Finally, since the decoder's message-update function is symmetric, Lemma 1 confirms the induction hypothesis in (8). ■

The last remaining step in establishing whether density evolution can be used with a decoder affected by WS deviations is to determine whether the error-correction performance of a code

concentrates around the cycle-free case. The property has been shown to hold in [9] (Theorems 2, 3 and 4) for an LDPC decoder affected by “wire noise” and “computation noise”. The wire noise model is similar to our deviation model, in the sense that the messages are passed through an additive noise channel, and that the noise applied to one message is independent of the noise applied to other messages. The proof presented in [9] only relies on the fact that the wire noise applied to a given message can only affect messages that are included in the directed neighborhood of the edge where it is applied, where the graph direction refers to the direction of message propagation. This clearly also holds in the case of our deviation model, and therefore the proof is the same.

Since the message error probability is independent of the transmitted codeword, and furthermore concentrates around the cycle-free case, density evolution can be used to determine the error-correction performance of a decoder perturbed by our deviation model, as long as the deviations are weakly symmetric.

C. Deviation Examples

As described in Section III-D, we collect deviation measurements from the test circuits by inputting test vectors representing random codewords, and distributed according to several $p_e^{(t-1)}$ values. We then generate estimates of the conditional distributions in (5). It is interesting to visualize the distributions using an aggregate measure such as the probability of observing a non-zero deviation

$$p_{\text{nz}}(\nu_{i,j}^{(t)}, x_i) = \mathbb{P}^{(p_e^{(t-1)})} \left(\mu_{i,j}^{(t)} \neq \nu_{i,j}^{(t)} \mid \nu_{i,j}^{(t)}, x_i \right). \quad (9)$$

These conditional probabilities are shown for a $(3, 30)$ circuit in Fig. 4. When $x_i = 1$, positive belief values indicate a correct decision, whereas when $x_i = -1$, negative belief values indicate a correct decision. We can see that in this example, deviations are more likely when the belief is incorrect than when it is correct, and therefore a symmetric deviation model is not consistent with these measurements. On the other hand, there is a sign symmetry between the “correct” part of the curves, and between the “incorrect” parts, that is $p_{\text{nz}}(\nu_{i,j}^{(t)}, 1) = p_{\text{nz}}(-\nu_{i,j}^{(t)}, -1)$, and for this reason a weakly symmetric model is consistent with the measurements. Note that the slight jaggedness observed for incorrect belief values of large magnitude in the $p_e^{(t-1)} = 0.008$ curves is due to the fact that these $\nu_{i,j}$ values occur only rarely. For the largest incorrect $\nu_{i,j}$




Fig. 4. Non-zero deviation probability given $\nu_{i,j}^{(t)}$ and x_i at two $p_e^{(t-1)}$ values, measured on a (3,30) circuit operated at $V_{dd} = 0.75$ V and $T_{clk} = 3.2$ ns. $3 \cdot 10^8$ decoding iteration trials were performed for each $p_e^{(t-1)}$ value. The total number of non-zero deviation events observed is 4,115,229 at $p_e^{(t-1)} = 0.015$, and 10,071,810 at $p_e^{(t-1)} = 0.008$.

values, only about 100 deviation events are observed for each point, despite the large number of Monte-Carlo (MC) trials.

Figure 5 shows a similar plot for a (3,6) circuit. In this case, $p_{nz}(\nu_{i,j}^{(t)}, x_i) \approx p_{nz}(-\nu_{i,j}^{(t)}, x_i)$, and a symmetric deviation model could be appropriate. Of course, since it is more general, a WS model is also appropriate.

Under the assumption that deviations are weakly symmetric, we have

$$\phi_{\mu_{i,j}^{(t)} | \nu_{i,j}^{(t)}, x_i}(\mu | \nu, 1) = \phi_{\mu_{i,j}^{(t)} | \nu_{i,j}^{(t)}, x_i}(-\mu | -\nu, -1).$$

Therefore, we can combine the $x_i = 1$ and $x_i = -1$ data to improve the accuracy of the estimated distributions.

Let p_L and p_H be respectively the smallest and largest $p_e^{(t-1)}$ values for which the deviations have been characterized. We can generate a conditional distribution for any $p_e^{(t-1)} \in [p_L, p_H]$ by interpolating from the nearest distributions that have been measured. We choose $p_H \geq p_e^{(0)}$ to make sure that the first iteration's deviation is within the characterized range. Because messages in the decoder are saturated once they reach the largest magnitude that can be represented, the circuit's switching activity decreases when the message error probability becomes very small. Since timing faults cannot occur when the circuit does not switch, we can expect deviations to be equally or less likely at $p_e^{(t-1)}$ values below p_L . Therefore, to define the deviation model




Fig. 5. Non-zero deviation probability given $\nu_{i,j}^{(t)}$ and x_i at two $p_e^{(t-1)}$ values, measured on a (3,6) circuit operated at $V_{dd} = 0.85$ V and $T_{clk} = 2.1$ ns. $3 \cdot 10^8$ decoding iteration trials were performed for each $p_e^{(t-1)}$ value. The total number of non-zero deviation events observed is 2,524,601 at $p_e^{(t-1)} = 0.09$, and 1,020,867 at $p_e^{(t-1)} = 0.05$.

for $p_e^{(t-1)} < p_L$, we make the pessimistic assumption that the deviation distribution remains the same as for $p_e^{(t-1)} = p_L$.

D. DE and Energy Curves

We evaluate the progress of the decoder affected by timing violations using quantized density evolution [30]. For the Offset Min-Sum algorithm, a DE iteration can be split into the following steps: 1-a) evaluating the distribution of the CN minimum, 1-b) evaluating the distribution of the CN output, after subtracting the offset, 2) evaluating the distribution of the ideal VN-to-CN message, and 3) evaluating the distribution of the faulty VN-to-CN messages. Step 1-a is given in [15], while the others are straightforward. In the context of DE, we write the message distribution as $\pi^{(t)} = \phi(\mu_{i,j}^{(t)} | x_i = 1)$, and the channel output distribution as $\pi^{(0)} = \phi(\mu_i^{(0)} | x_i = 1)$. We write a DE iteration as $\pi^{(t+1)} = f_\gamma(\pi^{(t)}, \pi^{(0)})$.

As mentioned in Section III-D, the energy consumption is modeled in terms of the message error probability and of the operating condition, and denoted $c_\gamma(p_e^{(t)})$. As for the deviation model, we use interpolation to define $c_\gamma(p_e^{(t)})$ for $p_e^{(t)} \in [p_L, p_H]$, and assume that $c_\gamma(p_e^{(t)}) = c_\gamma(p_L)$ for $p_e^{(t)} < p_L$. To display $f_\gamma(\pi^{(t)}, \pi^{(0)})$ and $c_\gamma(p_e^{(t)})$ on the same plot, we project $\pi^{(t)}$ onto the message error probability space.




Fig. 6. Examples of projected DE curves (solid lines) and energy curves (dashed lines) for rate 0.5 ensembles with $d_v \in \{3, 4, 5\}$, and $p_e^{(0)} = 0.09$.




Fig. 7. Examples of projected DE curves (solid lines) and energy curves (dashed lines) for the (3,30) and (4,40) ensembles (rate 0.9), with $p_e^{(0)} = 0.015$.

Several regular code ensembles were evaluated, with rates $\frac{1}{2}$ and $\frac{9}{10}$. Fig. 6 shows examples of projected DE curves and energy curves for rate- $\frac{1}{2}$ code ensembles with $d_v \in \{3, 4, 5\}$ and various operating conditions. The energy is measured as described in Appendix A and corresponds to one use of the test circuit (shown in Fig. 8). The nominal operating condition is $V_{dd} = 1.0$ V, $T_{clk} = 2.0$ ns and therefore these curves correspond to a reliable implementation. With a reliable

implementation, these ensembles have a channel threshold of $p_e^{(0)} \leq 0.12$ for the $(3, 6)$ ensemble, $p_e^{(0)} \leq 0.11$ for $(4, 8)$, and $p_e^{(0)} \leq 0.09$ for $(5, 10)$. We use $p_e^{(0)} = 0.09$ for all the curves shown in Fig. 6 to allow comparing the ensembles. As can be expected, a larger variable node degree results in faster convergence towards zero error rate, and it is natural to ask whether this property might provide greater fault tolerance and ultimately better energy efficiency. This is discussed in Section V-D.

Fig. 7 is a similar plot for the $(3, 30)$ and $(4, 40)$ ensembles. The channel threshold of both ensembles is approximately $p_e^{(0)} \leq 0.019$. For these curves, the nominal operating condition is $V_{dd} = 1.0$ V and $T_{clk} = 3$ ns. As we can see, the energy consumption per iteration of the $(4, 40)$ decoder is roughly double that of the $(3, 30)$ decoder. We note that in the case of the $(3, 30)$ ensemble, the reliable decoder stops making progress at an error probability of approximately 10^{-8} . This floor is the result of the message saturation limit chosen for the circuit.

V. ENERGY OPTIMIZATION

A. Design Parameters

As in a standard LDPC code-decoder design, the first parameter to be optimized is the choice of code ensemble. In this paper we restrict the discussion to regular codes, and therefore we need only to choose a degree pair (d_v, d_c) , where $R = 1 - d_v/d_c$ is the design rate of the code. For a fixed R , we can observe that both the energy consumption and the circuit area of the decoding circuit grow rapidly with d_v , and therefore it is only necessary to consider a few of the lowest d_v values.

Besides the choice of ensemble, we are interested in finding the optimal choice of operating parameters for the quasi-synchronous circuit. We consider here the supply voltage (V_{dd}) and the clock period (T_{clk}). Generally speaking, the supply voltage affects the energy consumption, while the clock period affects the decoding time, or latency. The energy and latency are also affected by the choice of code ensemble, since the number of operations to be performed depends on the node degrees. The operating parameters of a decoder are denoted as a vector $\gamma = [V_{dd}, T_{clk}]$.

The decoding of LDPC codes proceeds in an iterative fashion, and it is therefore possible to adjust the operating parameters on an iteration-by-iteration basis. In practice, this could be implemented in various ways, for example by using a pipelined sequence of decoder circuits, where each decoder is responsible for only a portion of the decoding iterations. It is also possible

to rapidly vary the clock frequency of a given circuit by using a digital clock divider circuit [31]. We denote by $\bar{\gamma}$ the sequence of parameters used at each iteration throughout the decoding, and we use $\bar{\gamma} = [\gamma_1^{N_1}, \gamma_2^{N_2}, \dots]$ to denote a specific sequence in which the parameter vector γ_1 is used for the first N_1 iterations, followed by γ_2 for the next N_2 iterations, and so on.

B. Objective

The performance of the LDPC code and of its decoder can be described by specifying a vector $\mathbf{P} = (p_e^{(0)}, p_{\text{res}}, T_{\text{dec}})$, where $p_e^{(0)}$ is the output error rate of the communication channel, p_{res} the residual error rate of VN-to-CN messages when the decoder terminates, and T_{dec} the expected decoding latency.

The decoder's performance \mathbf{P} and energy consumption E are controlled by $\bar{\gamma}$. The energy minimization problem can be stated as follows. Given a performance constraint $\mathbf{P} = (a, b, c)$, we wish to find the value of $\bar{\gamma}$ that minimizes E , subject to $p_e^{(0)} \geq a$, $p_{\text{res}} \leq b$, $T_{\text{dec}} \leq c$. As in the standard DE method, we propose to use the code's computation tree as a proxy for the entire decoder, and furthermore to use the energy consumption of the test circuit described in Appendix B as the optimization objective. To be able to replace the energy minimization of the complete decoder with the energy minimization of the test circuit, we make the following assumptions:

- 1) The ordering of the energy consumption is the same for the test circuit and for the complete decoder, that is, for any γ_1 and γ_2 , $E_{\text{TEST}}(\gamma_1) \leq E_{\text{TEST}}(\gamma_2)$ implies $E_{\text{DEC}}(\gamma_1) \leq E_{\text{DEC}}(\gamma_2)$, where $E_{\text{TEST}}(\gamma)$ and $E_{\text{DEC}}(\gamma)$ are respectively the energy consumption of the test circuit and of the complete decoder when using parameter γ .
- 2) The average message error rate in the test circuit and in the complete decoder is the same for all decoding iterations.
- 3) The latency of the complete decoder is proportional to the latency of the test circuit, that is, if $T_{\text{dec}}(\gamma)$ is the latency measured using the test circuit with parameter γ , the latency of the complete decoder is given by $\beta T_{\text{dec}}(\gamma)$, where β does not depend on γ .

Assumption 1 is reasonable because the test circuit is very similar to a computation unit used in the complete decoder. The difference between the two is that the test circuit only instantiates one full VNP, the remaining $(d_c - 1)$ VNPs being reduced to only their “front” part (as seen in Fig. 8), whereas the complete decoder has d_c full VNPs for every CNP. Assumption 2 is the

standard DE assumption, which is reasonable for sufficiently long codes. Finally, it is possible for the clock period to be slower in the complete decoder, because the increased area could result in longer interconnections between circuit blocks. Even if this is the case, the interconnect length only depends on the area of the complete decoder, which is not affected by the parameters we are optimizing, and hence β does not depend on γ .

Clearly, if Assumption 1 holds and the performance of the test circuit is the same as the performance of the complete decoder, then the solution of the energy minimization is also the same. The performance is composed of the three components $(p_e^{(0)}, p_{\text{res}}, T_{\text{dec}})$. The channel error rate $p_e^{(0)}$ does not depend on the decoder and is clearly the same in both cases. Because of Assumption 2, the complete decoder can achieve the same residual error rate as the test circuit when $p_e^{(0)}$ is the same. The latencies measured on the test circuit and on the complete decoder are not necessarily the same, but if Assumption 3 holds, and if we assume that the constant β is known, then we can find the solution to the energy minimization of the complete decoder subject to constraints $(p_e^{(0)}, p_{\text{res}}, T_{\text{dec}})$ by instead minimizing the energy of the test circuit with constraints $(p_e^{(0)}, p_{\text{res}}, T_{\text{dec}}/\beta)$.

We also consider another interesting optimization problem. It is well known that for a fixed degree of parallelism, energy consumption is proportional to processing speed (represented here by T_{dec}), which is observed both in the physical energy limit stemming from Heisenberg's uncertainty principle [32], as well as in practical CMOS circuits [33]. In situations where both throughput normalized to area and low energy consumption are desired, optimizing the product of energy and latency or *energy-delay product* (EDP) for a fixed circuit area can be a better objective. In that case the performance constraint is stated in terms of $\mathbf{P} = (p_e^{(0)}, p_{\text{res}})$, and the optimization problem becomes the following: given a performance constraint $\mathbf{P} = (a, b)$, minimize $E(\bar{\gamma}) \cdot T_{\text{dec}}(\bar{\gamma})$ subject to $p_e^{(0)} \geq a$, $p_{\text{res}} \leq b$, and a fixed circuit area.

C. Dynamic Programming

To solve the iteration-by-iteration energy and EDP minimization problems stated above, we adapt the ‘‘Gear-Shift’’ dynamic programming approach proposed in [20]. The original method relies on the fact that the message distribution has a 1-D characterization, which is chosen to be the error probability. By quantizing the error probability space, a trellis graph can be constructed in which each node is associated with a pair $(\tilde{p}_e^{(t)}, t)$. Quantized quantities are marked with tildes.

TABLE I
ENERGY AND EDP OPTIMIZATION RESULTS.

Code family	Nom. T_{clk}	Norm. area \dagger	$p_e^{(0)}$	p_{res}	Latency [ns]	Standard		Quasi-synchronous	
						Energy [pJ]	EDP [nJ · ns]	Best energy [pJ]	Best EDP [nJ · ns]
(3,6)	2.0 ns	1.066	0.12 ‡	$\leq 10^{-8}$	66	250	16.5	192 (-23%)	12.7 (-23%)
			0.09	$\leq 10^{-8}$	22	68.2	1.50	45.0 (-34%)	0.98 (-35%)
(4,8)	2.0 ns	1.44	0.09	$\leq 10^{-8}$	18	98.5	1.77	74.9 (-24%)	1.33 (-25%)
(3,30)	3.0 ns	1.099	0.019 ‡	$\leq 10^{-8}$	84.0	883	74.2	605 (-31%)	48.6 (-35%)
	2.5 ns	1.135	0.019 ‡	$\leq 10^{-8}$	70.0	916	64.1	664 (-28%)	46.5 (-27%)
	3.0 ns	1.099	0.015	$\leq 10^{-8}$	39.0	306	11.9	196 (-36%)	7.35 (-38%)
	2.5 ns	1.135	0.015	$\leq 10^{-8}$	32.5	324	10.5	214 (-34%)	6.92 (-34%)
(4,40)	3.0 ns	1.522	0.015	$\leq 10^{-8}$	27.0	364	9.83	224 (-38%)	5.93 (-40%)

\dagger Cell area divided by the minimal area of the smallest decoder having the same code rate. \ddagger Approx. threshold.

A particular choice of $\overline{\gamma}$ corresponds to a path P through the graph, and the optimization is transformed into finding the least expensive path that starts from the initial state $(\tilde{p}_e^{(0)}, 0)$ and reaches any state $(\tilde{p}_e^{(t)}, t)$ such that $\tilde{p}_e^{(t)} \leq p_{\text{res}}$ and the latency constraint is satisfied, if there is one. Note that to ensure that the solutions remain achievable in the original continuous space, the message error rates $p_e^{(t)}$ are quantized by rounding up. To maintain a good resolution at low error rates, we use a logarithmic quantization, with 1000 points per decade.

In the case of a faulty decoder, we want to evaluate the decoder's progress by tracking a complete message distribution using DE, rather than simply tracking the message error probability. In this case, the Gear-Shift method can be used as an approximate solver by projecting the message distribution $\boldsymbol{\pi}^{(t)} = \phi(\mu_{i,j}^{(t)} | x_i = 1)$ onto the error probability space. We refer to this method as DE-Gear-Shift. Any path through the graph is evaluated by performing DE on the entire path using exact distributions, but different paths are compared in the projection space. As a result, the solutions that are found are not guaranteed to be optimal, but they are guaranteed to accurately represent the progress of the decoder.

In the DE-Gear-Shift method, a path P is a sequence of states $\{\boldsymbol{\pi}^{(t)}\}$. As in the original Gear-Shift method, any sequence of decoder parameters $\overline{\gamma}$ corresponds to a path. We denote the projection of a state onto the error probability space as $p_e^{(t)} = \Theta(\boldsymbol{\pi}^{(t)})$. To each path P , we

associate an energy cost E_P and a latency cost T_P . A path ending at a state $\pi^{(t)}$ can be extended with one additional decoding iteration using parameter γ by evaluating one DE iteration to obtain $\pi^{(t+1)} = f_\gamma(\pi^{(t)}, \pi^{(0)})$. Performing this additional iteration adds an energy cost $c_\gamma(\tilde{p}_e^{(t)}, p_e^{(0)})$ and a latency cost T_γ to the path's cost. When optimizing EDP, we define the overall cost of a path C_P as $C_P = E_P \cdot T_P$. When optimizing energy under a latency constraint, we define the path cost as a two-dimensional vector $C_P = (E_P, T_P)$.

We use the following rules to discard paths that are suboptimal in the error probability space. Rule 1: Paths for which the message error rate is not monotonically decreasing are discarded. Rule 2: A path P with cost C_P is said to *dominate* another path P' with cost $C_{P'}$ if all the following conditions hold: 1) an ordering exists between C_P and $C_{P'}$, 2) $C_P \leq C_{P'}$, 3) $\Theta(\pi_P) \leq \Theta(\pi_{P'})$, where π_P denotes the last state reached by path P . The search for the least expensive path is performed breadth-first. After each traversal of the graph, any path that is dominated by another is discarded.

When the path cost is one-dimensional, the optimization requires evaluating $O(|\Gamma|N_s)$ DE iterations, where $|\Gamma|$ is the number of operating points being considered and N_s the number of quantization levels used for $\tilde{p}_e^{(t)}$. This can be seen from the fact that with a 1-D cost, Rule 2 implies that at most one path can reach a given state $\tilde{p}_e^{(t)}$. Therefore, $O(|\Gamma|N_s)$ DE iterations are required for each decoding iteration. In addition, upper bounds can be derived for the number of decoding iterations spanned by the trellis graph in terms of the smallest latency and energy cost of the parameters in Γ , and therefore it is a constant that does not depend on $|\Gamma|$ or N_s . On the other hand, when the cost is two-dimensional, the number of DE iterations could grow exponentially in terms of the number of decoding iterations. However, even in the case of a 2-D cost, an ordering exists between the costs of paths P and P' if $(E_P \geq E_{P'} \wedge T_P \geq T_{P'}) \vee (E_P \leq E_{P'} \wedge T_P \leq T_{P'})$, and in that case Rule 2 can be applied. In practice, for the cases presented in this paper, the discarding rules allowed to keep the number of paths down to a manageable level, even when using a 2-D cost. Note that an alternative to the use of a 2-D cost is to define a 1-D cost as $C_P = E_P + \kappa T_P$, and to perform a binary search for the value of κ that yields an optimal solution with the desired latency.

The algorithm can also be modified to search for parameter sequences that have other desirable properties beyond minimal energy or EDP. For example, if the decoder is implemented as a pipelined sequence of decoders, it can be desirable to favor solutions that do not require the

decoder to switch its parameters too often. We can find good approximate solutions by adding a penalty to E_P when the algorithm used in the current and next steps is different.

D. Results

We use DE-Gear-Shift to find good parameter sequences $\bar{\gamma}$ for several regular ensembles with rates $\frac{1}{2}$ and $\frac{9}{10}$. The parameter space Γ consists of (V_{dd}, T_{clk}) points with V_{dd} from 0.70 V to 1.0 V in steps of 0.05 V and several T_{clk} values depending on V_{dd} , in steps of 0.1 ns. The standard and quasi-synchronous decoders use the same circuits. Parameter α in (1) is set to $\alpha = 4$ for the (3, 6), (3, 30), and (4, 40) decoders, and to $\alpha = 2$ for the (4, 8) decoder. The offset parameter C in Alg. 1 is set to $C = 2$ for the (4, 40) decoder and to $C = 1$ for all other decoders. As part of our best effort to design a good standard circuit, in the case of the (3, 30) decoder we present results for two circuits synthesized with different nominal T_{clk} values. The standard circuit has a lower energy consumption when synthesized with $T_{clk} = 3$ ns, while it has a lower EDP when synthesized with $T_{clk} = 2.5$ ns.

We first run the DE-Gear-Shift solver without any path penalties to obtain the best possible parameter sequences, for both the energy and the EDP objectives. We also noticed that in some cases, adding a small algorithm change penalty allows to discover slightly better sequences. Note that when the objective is EDP, there is no constraint on latency. These results are summarized in Table I, where the energy is normalized per check node. Overall, we see that significant gains are possible while achieving the same channel noise, latency, and residual error requirements. The synthesis results show that increasing d_v while keeping the rate constant leads to a significant increase in circuit area. Despite this, increasing the node degrees can result in a reduction of the EDP. For the rate $\frac{9}{10}$ ensembles, going from $d_v = 3$ to $d_v = 4$ decreases EDP by 6.4% for a standard system, and by 14% for a quasi-synchronous system. However this is not the case for the rate $\frac{1}{2}$ ensembles, where $d_v = 3$ has the smaller EDP. As expected, we can also see that much more energy is required when the channel quality is close to the ensemble's threshold.

By applying a cost penalty to parameter switches, it is possible to find parameter sequences with few switches, without a large increase in cost. For example, for a (3, 6) decoder starting at $p_e^{(0)} = 0.09$, a single operating condition can provide a 32% EDP improvement, using $\bar{\gamma} = [[0.8 \text{ V}, 2.1 \text{ ns}]^{11}]$. The probability of a non-zero deviation in that schedule ranges from 0.6% to 7.2%. In the case of a (3, 30) decoder synthesized at a nominal $T_{clk} = 2.5$ ns, for $p_e^{(0)} = 0.015$

the sequence $\bar{\gamma} = [[0.8 \text{ V}, 2.5 \text{ ns}]^{12}, [1.0 \text{ V}, 2.5 \text{ ns}]]$ provides a 30% EDP improvement, with non-zero deviation probabilities from 0 to 0.8%. For a (4, 40) decoder, the single-parameter sequence $\bar{\gamma} = [[0.8 \text{ V}, 2.8 \text{ ns}]^9]$ provides a 39% EDP improvement, with non-zero deviation probabilities from 1.6 to 4.6%.

VI. CONCLUSION

We presented a method for the design of synchronous circuit implementations of signal processing algorithms that permits timing violations without the need for hardware compensation. We introduced a model for the deviations occurring in LDPC decoder circuits affected by timing faults that represent the circuit behavior accurately [27], while being independent of the iterative progress of the decoder. In addition, we showed that in order to use density evolution to predict the performance of the faulty decoder, it is sufficient for the deviation model to have a weak symmetry property, which is more general than previously proposed sufficient properties.

We then presented an approximate optimization method called DE-Gear-Shift to find sequences of circuit operating parameters that minimize the energy or the energy-delay product. The method is similar to the previously proposed Gear-Shift method, but relies on density evolution rather than ExIT charts to evaluate the average iterative progress of the decoder. Our results show that the best energy or EDP reduction is achieved by operating the circuit with a large number of timing violations (often with an average probability of non-zero deviation above 1%). Furthermore, important savings can be achieved with few parameter switches, and without any compromise on circuit area or decoding performance.

In this work, we only considered delay variations associated with the signal transitions at the input of the circuit. While the energy savings that result from tolerating these variations are already significant, we ultimately see quasi-synchronous systems as an approach for tolerating the large process variations found in near-threshold CMOS circuits and other emerging computing technologies, potentially enabling energy savings of an order of magnitude. Furthermore, we believe this approach can be extended to other self-correcting algorithms, such as deep neural networks.

APPENDIX A

CAD WORKFLOW

The deviations and the energy consumption are measured directly on optimized circuit models generated by a commercial synthesis tool (Cadence Encounter [34]). We use TSMC’s 65 nm process with the *tcbn65gplus* cell library [35]. In order to provide a fair assessment of the improvements provided by the quasi-synchronous circuit, we first synthesize a *benchmark* circuit that represents a best effort at optimizing the metric of interest, for example energy consumption. Since we do not have a specific throughput constraint for the design, we synthesize the benchmark circuit at the standard supply voltage of the library ($V_{dd} = 1.0V$), while the clock period is chosen as small as possible without causing a degradation of the target metric. Second, we synthesize a *nominal* circuit that will serve as the basis for the quasi-synchronous design. In this work, we use a standard synthesis algorithm for the nominal circuit, and in all the cases that we report on, the nominal and the benchmark circuits are actually the same. Using a standard synthesis method for the nominal circuit allows using off-the-shelf tools, but is not ideal since the objective of a standard synthesis algorithm (to make all paths only as fast as the clock period) differs from the objective pursued when some timing violations are permitted. For example, results in [36] show that the power consumption of a circuit can be reduced by up to 32% when the gate-sizing optimization takes into account the acceptable rate of timing violations. Therefore it is possible that our results could be improved by using a different synthesis algorithm.

Once the circuit is synthesized, we perform a static timing analysis of the gate-level model at various supply voltages. All timing analyses (including at the nominal supply) are performed using timing libraries generated by the Cadence Encounter Library Characterization tool. We then use this timing information in a functional simulation of the gate-level circuit to observe the dynamic effect of path delay variations and measure the deviation statistics. Any source of delay variation that can be simulated can be studied, but in this paper we focus on variations due to path activation, that is the variations in delay caused by the different propagation times required by different input transitions. Note that other methods could be used to obtain the propagation delays, such as the method described in [37] based on analytical models. In addition to speeding up the characterization, such methods allow considering the effect of process variations.

Power estimation is performed by collecting switching activity data in the functional simulation

and using the power estimation engine in Cadence Encounter. However, because the circuit is operated in a quasi-synchronous manner, the clock period used to run the circuit is not necessarily the same as the nominal clock period. When that is the case, the power estimation generated by the synthesis tool cannot be used directly. First, the switching activity recorded during the functional simulation must be scaled so that it corresponds to the nominal clock period. The tool's power estimation then reports the dynamic power P_{dyn} and the static power P_{stat} . The dynamic energy consumed during one clock cycle does not depend on the clock period, whereas the static energy does. Therefore, the total energy consumed during one cycle by the quasi-synchronous circuit is given by $E_{\text{cycle}} = P_{\text{dyn}}T_{\text{clk,nom}} + P_{\text{stat}}T_{\text{clk}}$, where $T_{\text{clk,nom}}$ is the nominal clock period and T_{clk} is the actual clock period used to run the circuit.

APPENDIX B

TEST CIRCUIT MONTE-CARLO SIMULATION

A suitable test circuit for a row-layered decoder architecture consists in implementing a single check node processor, as well as the necessary logic taken from the variable node processor block to send d_v messages to the CNP, and receive one message from the CNP. This test circuit is shown in Fig. 8. It re-uses logic blocks that are found in the complete decoder, ensuring the accuracy of the deviation and energy measurements, and minimizing design time.

The test circuit is used to evaluate the decoder's computation tree (shown in Fig. 3). The VNP with index 1, shown at the top, is always mapped to the VN that is at the head of the computation tree, while the VNPs at the bottom of the figure are mapped to different VNs as the CNP is successively mapped to each CN neighbor of the head VN.

At any given clock cycle, a *VNP front* block is mapped to a particular VN i . For illustrative purposes, we simply index the VN neighbors from 1 to d_c , even if the VNs mapped to the bottom VNPs actually change at each layer. Each *VNP front* block takes as input the previous belief total of that VN Λ'_i , and the previous CN-to-VN message corresponding to layer ℓ , $\lambda_{i,J(i,\ell)}^{(t-1)}$.

To perform the Monte-Carlo simulation, a *VNP front* circuit block with index i must send a message $\mu_i^{(t)}$, randomly generated according to a 1-D normal distribution with error probability $p_e^{(t)}$. However, the only inputs that are controllable are Λ'_i and $\lambda_{i,J(i,\ell)}^{(t-1)}$. To simplify the Monte-Carlo simulation, we disregard the true distribution of $\lambda_{i,J(i,\ell)}^{(t-1)}$ and generate it according to a 1-D normal distribution. We also introduce another simplification: we assume that messages received

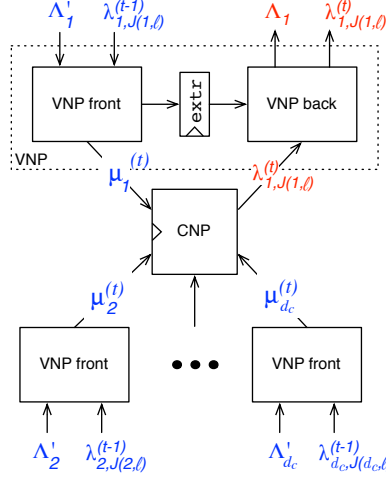


Fig. 8. Block diagram of the test circuit.

at a VN only modify the total belief at the end of the iteration, as would be the case when using a flooding schedule. As a result, the messages $\mu_i^{(t)}$ are identically distributed with error rate parameter $p_e^{(t)}$ for all ℓ . Note that these simplifications are not necessary, and they could be removed at the cost of a slightly more cumbersome Monte-Carlo simulation.

To generate inputs with the appropriate distribution, we use the fact that $\Lambda'_i = \mu_{i,J(i,\ell)}^{(t)} + \lambda_{i,J(i,\ell)}^{(t-1)}$. On a cycle-free Tanner graph, $\mu_{i,J(i,\ell)}^{(t)}$ and $\lambda_{i,J(i,\ell)}^{(t-1)}$ are independent, but naturally Λ'_i and $\lambda_{i,J(i,\ell)}^{(t-1)}$ are not. Therefore, we generate $\mu_{i,J(i,\ell)}^{(t)}$ and $\lambda_{i,J(i,\ell)}^{(t-1)}$ and sum them to obtain Λ'_i .

To complete the DE iteration, we want to measure an extrinsic message belonging to the next iteration. Because we assume a flooding schedule, this extrinsic message can be obtained by summing any set of $(d_v - 1)$ messages in the current iteration. To achieve this, we start a DE iteration by setting $\Lambda'_1 \leftarrow 0$ and $\lambda_{1,J(1,\ell)}^{(t-1)} \leftarrow 0$ for all ℓ . The desired extrinsic message then corresponds to the total belief output of the circuit $\Lambda_1^{(t)}$ after $d_v - 1$ layers have been evaluated.

Just like the processor used in the complete decoder, the test circuit has one input and one output register, as well as one internal pipeline register, for a latency of 3 clock cycles. In order to keep the pipeline fed, several distinct computation trees are evaluated in parallel during the Monte-Carlo simulation.

ACKNOWLEDGEMENTS

The authors wish to thank CMC Microsystems for providing access to the Cadence tools and TSMC 65nm CMOS technology, and Gilles Rust for advice on Cadence tools and cell library characterization.

REFERENCES

- [1] S. Ghosh and K. Roy, "Parameter variation tolerance and error resiliency: New design paradigm for the nanoscale era," *Proc. of the IEEE*, vol. 98, no. 10, pp. 1718–1751, Oct. 2010.
- [2] R. Dreslinski, M. Wieckowski, D. Blaauw, D. Sylvester, and T. Mudge, "Near-threshold computing: Reclaiming Moore's law through energy efficient integrated circuits," *Proc. of the IEEE*, vol. 98, no. 2, pp. 253–266, Feb. 2010.
- [3] T. Austin, V. Bertacco, D. Blaauw, and T. Mudge, "Opportunities and challenges for better than worst-case design," in *Proc. of the 2005 Asia and South Pacific Design Automation Conf.*, 2005, pp. 2–7.
- [4] K. Bowman, J. Tschanz, N. S. Kim, J. Lee, C. Wilkerson, S.-L. Lu, T. Karnik, and V. De, "Energy-efficient and metastability-immune resilient circuits for dynamic variation tolerance," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 49–63, Jan. 2009.
- [5] S. Das, C. Tokunaga, S. Pant, W.-H. Ma, S. Kalaiselvan, K. Lai, D. Bull, and D. Blaauw, "RazorII: In situ error detection and correction for PVT and SER tolerance," *IEEE J. of Solid-State Circuits*, vol. 44, no. 1, pp. 32–48, Jan. 2009.
- [6] R. Hegde and N. R. Shanbhag, "Energy-efficient signal processing via algorithmic noise-tolerance," in *Proc. 1999 Int. Symp. on Low Power Electronics and Design*, Aug. 1999, pp. 30–35.
- [7] B. Shim, S. Sridhara, and N. Shanbhag, "Reliable low-power digital signal processing via reduced precision redundancy," *IEEE Trans. on VLSI Systems*, vol. 12, no. 5, pp. 497–510, May 2004.
- [8] Y. Liu, T. Zhang, and K. Parhi, "Computation error analysis in digital signal processing systems with overscaled supply voltage," *IEEE Trans. on VLSI Systems*, vol. 18, no. 4, pp. 517–526, Apr. 2010.
- [9] L. Varshney, "Performance of LDPC codes under faulty iterative decoding," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4427–4444, Jul. 2011.
- [10] F. Leduc-Primeau and W. J. Gross, "Faulty Gallager-B decoding with optimal message repetition," in *Proc. 50th Allerton Conf. on Communication, Control, and Computing*, Oct. 2012.
- [11] S. M. S. Tabatabaei Yazdi, H. Cho, and L. Dolecek, "Gallager B decoder on noisy hardware," *IEEE Trans. Commun.*, vol. 61, no. 5, pp. 1660–1673, May 2013.
- [12] C.-H. Huang, Y. Li, and L. Dolecek, "Gallager B LDPC decoder with transient and permanent errors," *IEEE Trans. Commun.*, vol. 62, no. 1, pp. 15–28, Jan. 2014.
- [13] C. Ngassa, V. Savin, and D. Declercq, "Min-sum-based decoders running on noisy hardware," in *IEEE Global Communications Conf. (GLOBECOM)*, Dec. 2013, pp. 1879–1884.
- [14] E. Dupraz, D. Declercq, B. Vasic, and V. Savin, "Analysis and design of finite alphabet iterative decoders robust to faulty hardware," *IEEE Trans. on Communications*, vol. 63, no. 8, pp. 2797–2809, Aug. 2015.
- [15] A. Balatsoukas-Stimming and A. Burg, "Density evolution for min-sum decoding of LDPC codes under unreliable message storage," *IEEE Commun. Lett.*, vol. 18, no. 5, pp. 849–852, May 2014.
- [16] S. Brkic, O. Al Rasheed, P. Ivanis, and B. Vasic, "On fault tolerance of the Gallager B decoder under data-dependent gate failures," *IEEE Commun. Lett.*, vol. 19, no. 8, pp. 1299–1302, Aug. 2015.

- [17] C. G. Blake and F. R. Kschischang, "Energy consumption of VLSI decoders," *IEEE Trans. on Information Theory*, vol. 61, no. 6, pp. 3185–3198, Jun. 2015.
- [18] C. Blake and F. R. Kschischang, "On the energy complexity of LDPC decoder circuits," *CoRR*, vol. abs/1502.07999, 2015. [Online]. Available: <http://arxiv.org/abs/1502.07999>
- [19] K. Ganesan, P. Grover, J. Rabaey, and A. Goldsmith, "On the total power capacity of regular-LDPC codes with iterative message-passing decoders," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 2, pp. 375–396, Feb. 2016.
- [20] M. Ardakani and F. R. Kschischang, "Gear-shift decoding," *IEEE Trans. Commun.*, vol. 54, no. 7, pp. 1235–1242, Jul. 2006.
- [21] K. Cushon, C. Leroux, S. Hemati, S. Mannor, and W. Gross, "A min-sum iterative decoder based on pulsewidth message encoding," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 57, no. 11, pp. 893–897, Nov. 2010.
- [22] C. Roth, P. Meinerzhagen, C. Studer, and A. Burg, "A 15.8 pJ/bit/iter quasi-cyclic LDPC decoder for IEEE 802.11n in 90 nm CMOS," in *2010 IEEE Asian Solid State Circuits Conf.*, Nov. 2010, pp. 1–4.
- [23] Y. Sun, G. Wang, and J. Cavallaro, "Multi-layer parallel decoding algorithm and VLSI architecture for quasi-cyclic LDPC codes," in *IEEE Int. Symp. on Circuits and Systems*, May 2011, pp. 1776–1779.
- [24] A. Cevrero, Y. Leblebici, P. Ienne, and A. Burg, "A 5.35 mm² 10GBASE-T ethernet LDPC decoder chip in 90nm CMOS," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2010.
- [25] C.-L. Wey, M.-D. Shieh, and S.-Y. Lin, "Algorithms of finding the first two minimum values and their hardware implementation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 11, pp. 3430–3437, Dec. 2008.
- [26] B. Sedighi, N. P. Anthapadmanabhan, and D. Suvakovic, "Timing errors in LDPC decoding computations with overscaled supply voltage," in *Proc. 2014 Int. Symp. on Low Power Electronics and Design*, Aug. 2014, pp. 201–206.
- [27] F. Leduc-Primeau and W. J. Gross, "Finite-length quasi-synchronous LDPC decoders," in *9th Int. Symp. on Turbo Codes and Iterative Information Processing*, Sep. 2016, pp. 325–329.
- [28] T. J. Richardson and R. L. Urbanke, "The capacity of low-density parity-check codes under message-passing decoding," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 599–618, Feb. 2001.
- [29] T. Richardson and R. Urbanke, *Modern Coding Theory*. Cambridge University Press, 2008.
- [30] S.-Y. Chung, J. David Forney, T. J. Richardson, and R. Urbanke, "On the design of low-density parity-check codes within 0.0045dB of the Shannon limit," *IEEE Commun. Lett.*, vol. 5, no. 2, pp. 58–60, Feb. 2001.
- [31] T. Fischer, J. Desai, B. Doyle, S. Naffziger, and B. Patella, "A 90-nm variable frequency clock system for a power-managed itanium architecture processor," *IEEE J. Solid-State Circuits*, vol. 41, no. 1, pp. 218–228, Jan. 2006.
- [32] S. Lloyd, "Ultimate physical limits to computation," *Nature*, vol. 406, no. 6799, pp. 1047–1054, Aug. 2000.
- [33] R. Gonzalez and M. Horowitz, "Energy dissipation in general purpose microprocessors," *IEEE J. Solid-State Circuits*, vol. 31, no. 9, pp. 1277–1284, Sep. 1996.
- [34] Cadence Design Systems Inc., "Encounter digital implementation system." [Online]. Available: http://www.cadence.com/products/di/edi_system
- [35] *TCBN65GPLUS TSMC 65nm Core Library Databook*, Taiwan Semiconductor Manufacturing Company, Ltd.
- [36] A. Kahng, S. Kang, R. Kumar, and J. Sartori, "Slack redistribution for graceful degradation under voltage overscaling," in *Proc. 15th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2010, pp. 825–831.
- [37] A. Pirbadian, M. Khairy, A. Eltawil, and F. Kurdahi, "State dependent statistical timing model for voltage scaled circuits," in *IEEE Intl Symp. on Circuits and Systems (ISCAS)*, Jun. 2014, pp. 1432–1435.